HP-UX

HP StorageWorks Disk Array XP operating system configuration guide

XP48
XP128
XP512
XP1024
XP12000

fifth edition (August 2004)

part number: A5951-96014

This guide describes the requirements and procedures for connecting the XP family of disk arrays to an HP-UX system and configuring the new disk array for operation with HP-UX.

Printed in the U.S.A.

HP StorageWorks Disk Array XP Operating System Configuration Guide: HP-UX

fifth edition (Auguust 2004)
part number: A5951-96014

Contents

About this guide

This guide describes the requirements and procedures for connecting the XP family of disk arrays to an HP-UX system, and configuring the new disk array for operation with HP-UX.

Intended audience

This guide is intended for system administrators who have knowledge of the following topics:

- Data processing concepts
- Direct access storage device subsystems and their basic functions
- Disk arrays and RAID technology
- Operating system commands and utilities

Disk arrays

Unless otherwise noted, the term *disk array* refers to these disk arrays:

HP Surestore Disk Array XP512
HP Surestore Disk Array XP48
HP StorageWorks Disk Array XP128
HP StorageWorks Disk Array XP1024
HP StorageWorks XP12000 Disk Array

Related documentation

HP provides the following related documentation:

- *HP StorageWorks Disk Array XP128: Owner's Guide*
- *HP StorageWorks Disk Array XP1024: Owner's Guide*
- *HP StorageWorks XP12000 Disk Array: Owner's Guide*

For information about operating system commands and third-party products, refer to the manufacturer's documentation.

Conventions

This guide uses the following text conventions.

Figure 1	Blue text represents a cross-reference. For the online version of this guide, the reference is linked to the target.
www.hp.com	Underlined, blue text represents a website on the Internet. For the online version of this guide, the reference is linked to the target.
literal	Bold text represents literal values that you type exactly as shown, as well as key and field names, menu items, buttons, file names, application names, and dialog box titles.
variable	Italic type indicates that you must supply a value. Italic type is also used for manual titles.
`input/output`	Monospace font denotes user input and system responses, such as output and messages.
Example	Denotes an example of input or output. The display shown in this guide may not match your configuration exactly.
[ ]	Indicates an optional parameter.
{ }	Indicates that you must specify at least one of the listed options.
|	Separates alternatives in a list of options.

HP technical support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support:

http://h18006.www1.hp.com/storage/arraysystems.html

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP storage website

For the most current information about HP StorageWorks XP products, visit the support website. Select the appropriate product or solution from this website:

http://h18006.www1.hp.com/storage/arraysystems.html

For information about product availability, configuration, and connectivity, consult your HP account representative.

HP authorized reseller

For the name of your nearest HP authorized reseller, you can obtain information by telephone:

United States	1-800-345-1518
Canada	1-800-263-5868
Or contact:	www.hp.com

Revision history

September, 1999	Open-8 emulation added.
January, 2000	Content extensively revised and reorganized.
June, 2000	Added support for XP512. Content reorganized and revised.
February, 2001	Added appendixes C, D, E, and F. Added glossary.
April 2003	Changed brand name to StorageWorks. Added support for OPEN-L and OPEN-V. Changed CVS to VCS.
November 2003	Revised I/O timeout procedure. Minor update.
August 2004	Updated for XP12000 and general improvements.

Warranty statement

HP warrants that for a period of ninety calendar days from the date of purchase, as evidenced by a copy of the invoice, the media on which the Software is furnished (if any) will be free of defects in materials and workmanship under normal use.

DISCLAIMER. EXCEPT FOR THE FOREGOING AND TO THE EXTENT ALLOWED BY LOCAL LAW, THIS SOFTWARE IS PROVIDED TO YOU "AS IS" WITHOUT WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED. HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow exclusions of implied warranties or conditions, so the above exclusion may not apply to you to the extent prohibited by such local laws. You may have other rights that vary from country to country, state to state, or province to province.

WARNING! **YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT USE OF THE SOFTWARE IS AT YOUR SOLE RISK.** HP DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, VIRUS-FREE OR ERROR-FREE, OR THAT DEFECTS IN THE SOFTWARE WILL BE CORRECTED. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. HP DOES NOT WARRANT OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE OR RELATED DOCUMENTATION IN TERMS OF THEIR CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY HP OR HP'S AUTHORIZED REPRESENTATIVES SHALL CREATE A WARRANTY.

LIMITATION OF LIABILITY. EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT INCLUDING NEGLIGENCE WILL HP OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Your use of the Software is entirely at your own risk. Should the Software prove defective, you assume the entire cost of all service, repair or correction. Some jurisdictions do not allow the exclusion or limitation of liability for incidental or consequential damages, so the above limitation may not apply to you to the extent prohibited by such local laws.

NOTE. EXCEPT TO THE EXTENT ALLOWED BY LOCAL LAW, THESE WARRANTY TERMS DO NOT EXCLUDE, RESTRICT OR MODIFY, AND ARE IN ADDITION TO, THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE LICENSE OF THE SOFTWARE TO YOU; PROVIDED, HOWEVER, THAT THE CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS IS SPECIFICALLY DISCLAIMED AND SHALL NOT GOVERN OR APPLY TO THE SOFTWARE PROVIDED IN CONNECTION WITH THIS WARRANTY STATEMENT.

1

Installation

Installation of the HP StorageWorks Disk Array XP is performed by your HP service representative and you. The HP service representative installs the disk array and formats the disk devices. You configure the host server for the new devices with assistance from the HP service representative.

Features and requirements

The disk array and host have the following features and requirements.

- HP StorageWorks disk arrays:

  XP48: Up to 48 drives from 72 GB to 8.7 TB, 24 FC ports

  XP128: From 8 to 128 drives for up to 18 TB, 48 FC ports

  XP512: Up to 512 drives from 72 GB to 93 TB, 48 FC ports

  XP1024: From 8 to 1024 drives for up to 149 TB, 64 FC ports

  XP12000: Up to 1152 drives for up to 165 TB, 128 FC ports
- HP-UX-supported processor
- HP-UX, version 10.20, 11.0, or 11i v2 with the latest patches (if any)
- Host Bus Adapters (HBAs): Install adapters and all utilities and drivers. Refer to the adapter documentation for installation details.
- *(Recommended)* HP StorageWorks Command View XP with LUN management feature or Remote Control with the LUN Configuration Manager XP option for configuring disk array ports and paths.
- *(Recommended)* HP StorageWorks Secure Manager XP: Allows the host to access only array devices for which it is authorized.
- Other available XP Software (some may not apply to your system):
  HP StorageWorks Business Copy XP
  HP StorageWorks Continuous Access XP
  HP StorageWorks Continuous Access Extension XP
  HP StorageWorks Auto LUN XP
  HP StorageWorks Data Exchange XP
  HP StorageWorks Resource Manager XP
  HP StorageWorks RAID Manager XP
  HP StorageWorks Cache LUN XP
  HP StorageWorks Auto Path XP
  HP StorageWorks Cluster Extension XP
  HP StorageWorks Performance Advisor XP software

Fibre Channel interface

The XP48, XP128, XP512, XP1024, and XP12000 disk arrays support these 1 Gbps and 2 Gbps Fibre Channel interfaces:

- Short-wave non-OFC (open fiber control) optical interface
- Multimode optical cables with SC or LC connectors
- Public or private arbitrated loop (FC-AL) or fabric direct attach
- Fibre Channel switches

Even though the interface is Fibre Channel, this guide uses the term "SCSI disk" because disk array devices are defined to the host as SCSI disks.

Device emulation types

The disk arrays support the following device emulation types:

- **OPEN-K/3/8/9/E devices:** OPEN-x logical units represent disk devices. Except for OPEN-V, these devices are based on fixed sizes. OPEN-V is a user-defined size. Supported emulations include OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L, and OPEN-V devices.
- **LUSE devices (OPEN-*x***n*):** Logical Unit Size Expansion (LUSE) allows you to combine 2 to 36 OPEN-x devices to create expanded LDEVs larger than standard OPEN-K/3/8/9/E disk devices. For example, an OPEN-x LUSE volume created from ten OPEN-x CVS volumes is designated as OPEN-x*10.
- **CVS devices (OPEN-x CVS):** Volume Size Configuration (VSC) defines custom volumes (CVS) that are smaller than normal fixed-sized logical disk devices (volumes). (OPEN-V is a CVS-based custom disk size that you determine. OPEN-L does not support CVS.)
- **LUSE (expanded) CVS devices (OPEN-x*n CVS):** LUSE CVS combines CVS devices to create an expanded device. This is done by first creating CVS custom-sized devices and then using LUSE to combine from 2 to 36 CVS devices. For example, if three OPEN-9 CVS volumes are combined to create an expanded device, this device is designated as OPEN-9*3-CVS.

Failover

The disk array supports standard products that provide host and/or application failover, I/O path failover, and logical volume management (LVM). For HP-UX, the disk array supports the following:

- HP Multi-Computer/ServiceGuard (MC/ServiceGuard) software for application failover
- Alternate link for I/O path failover (included in HP-UX)
- Logical volume management (included in HP-UX)

SNMP configuration

The disk arrays support standard Simple Network Management Protocol (SNMP) for remotely managing the disk array from the host. The SNMP agent on the remote console PC or Command View can provide status and Remote Service Information Message (R-SIM) reporting to the SNMP manager on the host for up to eight disk arrays. To configure the SNMP manager on the host, refer to the operating system documentation.

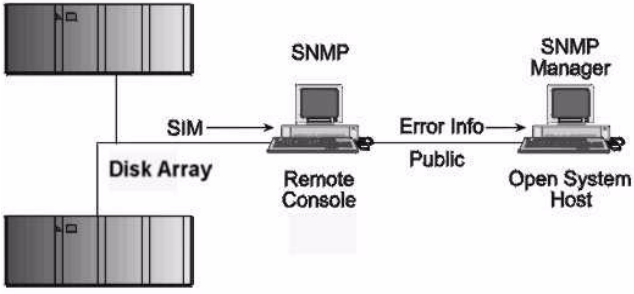

RAID Manager command devices

RAID Manager manages Business Copy (BC) and/or Continuous Access (CA) operations from a server host. To use RAID Manager with BC or CA, you must use Command View or LUN Configuration Manager to designate at least one LDEV as a command device. Refer to the Command View or LUN Configuration Manager user guide for information about how to designate a command device.

Installation procedures

The HP representative and you perform the following procedures.

1. “Install and configure the disk array” (page 16).

   “Setting the Host Mode for the disk array ports”

   “Setting the System Option Modes”

   “Configuring the Fibre Channel ports”
2. “Install and configure the host” (page 20).

   “Loading the OS and software”

   “Installing and configuring the HBAs”

   “Clustering and Fabric zoning”

   “Fabric zoning and LUN security for multiple operating systems”
3. “Connect the disk array” (page 23).

   “Defining the paths”

   “Verifying HBA installation”

   “Verifying device recognition”
4. “Configure disk array devices” (page 28).

   “Verifying the device files and drivers”

   “Creating the device files”

   “Creating the physical volumes”

   “Creating new volume groups”

   “Creating logical volumes”

   “Creating the file systems”

   “Setting the I/O timeout parameter”

   “Creating the mount directories”

   “Mounting and verifying the file systems”

   “Setting and verifying the auto-mount parameters”

Install and configure the disk array

The HP service representative performs the following tasks:

- Assembling hardware and installing software
- Loading the microcode updates
- Installing the channel adapters (CHAs) and cabling
- Installing and formatting devices

You perform the additional tasks below. If you do not have Command View or LUN Configuration Manager, your HP service representative can perform these tasks for you.

Setting the Host Mode for the disk array ports

The disk array ports have Host Modes that you must set depending on the host you use. After the disk array is installed, use Command View (shown) or LUN Configuration Manager to set the Host Mode for each port.

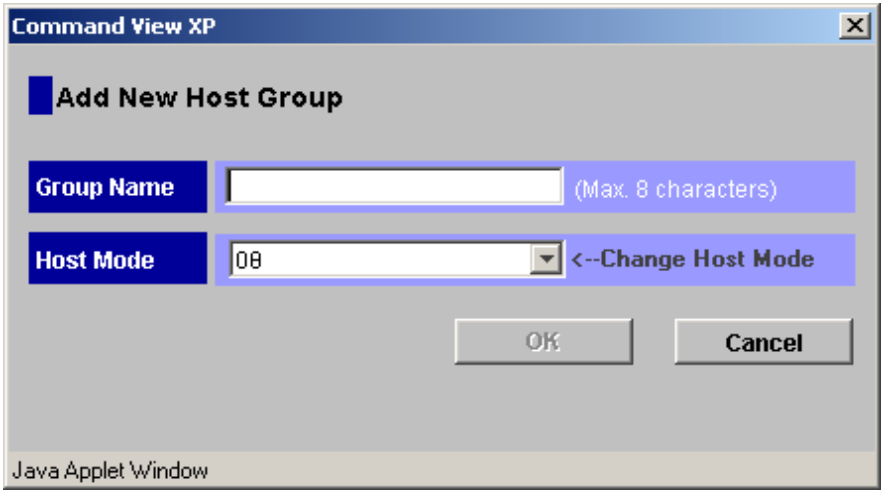

The host mode setting for HP-UX is 08.

Setting the System Option Modes

The HP representative sets the System Option Mode(s) based on the operating system and software configuration of the host.

Mode	Level	HA Software	Description and Usage
140	Optional	None	ON: Response to the Inquiry command is changed, and the volume can be used from VeritasNetBackUP in heterogeneous OS configuration, such as HP-UX. OFF: Normal response to the Inquiry command.
186	Mandatory	Veritas Database Editions/ Advanced Cluster	ON: Mandatory setting when VERITAS Database Editions/Advanced Cluster is connected. OFF: VERITAS Database Editions/Advanced Cluster should not be connected with this setting.
254	Optional	Veritas Database Editions	Change the response of reserve conflict status to Read Capacity, Verify, and Start Stop Unit. ON: Normal end is reported. OFF: Reserve conflict status is reported.
280	Optional		HP-UX Ghost LUN Remove OS version: HP-UX 11.0 and later. ON: When the host scans the LUNS on the port, the disk array is not represented in the device list for HP-UX. OFF: When the host scans the LUNS on the port, the disk array is represented in the device list for HP-UX.

Configuring the Fibre Channel ports

Configure the disk array Fibre Channel ports by using Command View (shown) or the Fibre Parameter window in LUN Configuration Manager. Select the settings for each port based on your storage area network topology. Use switch zoning if you connect different types of hosts to the array through the same switch.

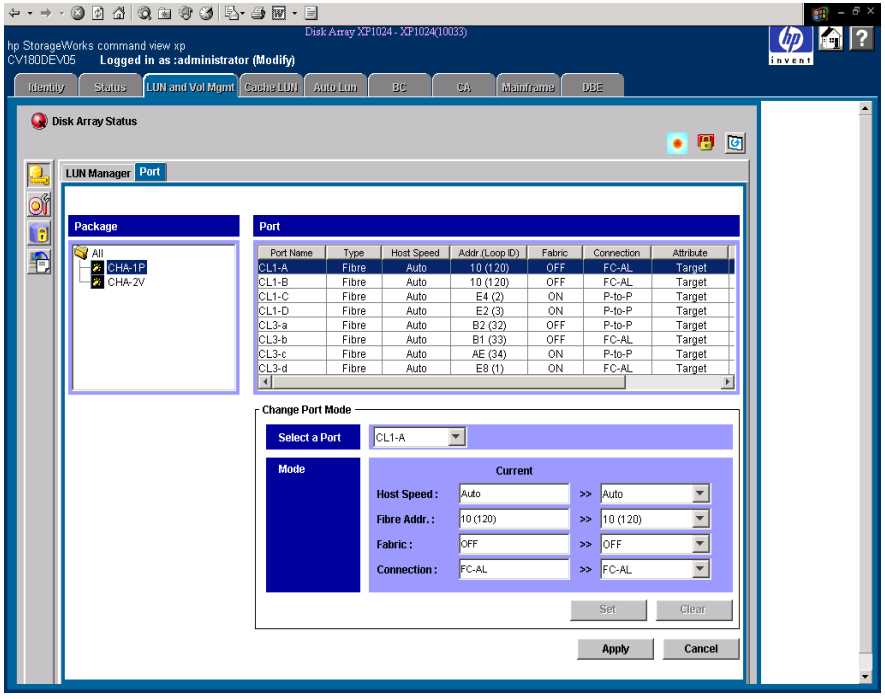

Fibre Address

In fabric environments, the port addresses are assigned automatically. In arbitrated loop environments, you set the port addresses by selecting a unique arbitrated loop physical address (AL-PA) or loop ID for each port.

Fabric and Connection parameter settings

You can set each array port to FABRIC ON or OFF with connections of POINT-TO-POINT or FC-AL as shown in the following table and figures. For detailed topology information, refer to the *HP StorageWorks SAN Design Reference Guide* on the hp.com website.

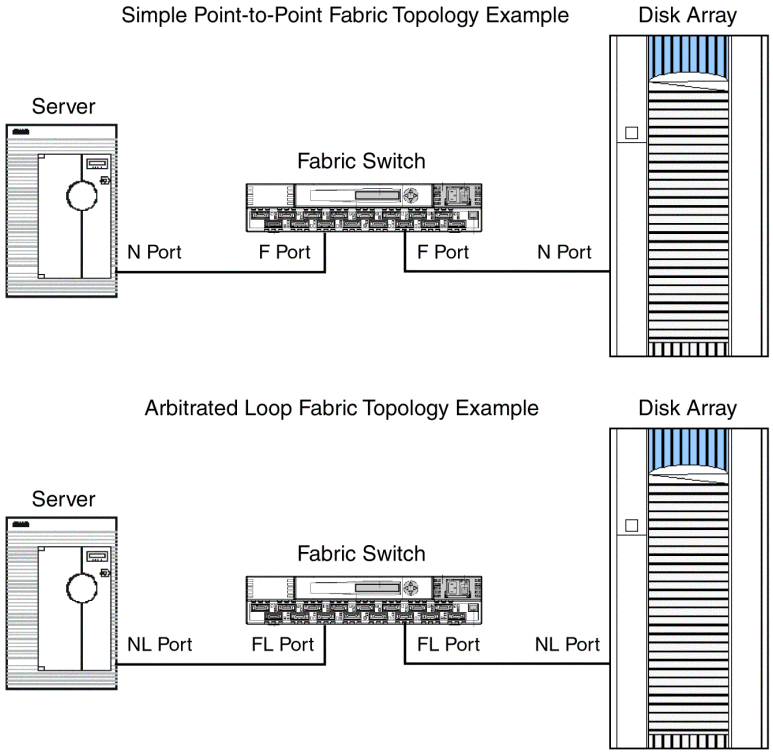

Table 1. Fibre Topology Settings on the Remote Console

Fabric parameter	Connection parameter	Provides
ON	FC-AL	Not supported
ON	Direct Fabric Attach	F-port (fabric port)
OFF	FC-AL	AL-port (private arbitrated loop)
OFF	Direct Fabric Attach	Not supported

Install and configure the host

Install and configure the host and host bus adapters (HBAs) that connect the host to the disk array.

Loading the OS and software

Follow the manufacturer's instructions to load the operating system and software onto the host. Load all OS patches and configuration utilities supported by HP and the HBA manufacturer.

Installing and configuring the HBAs

Install and configure the host bus adapters using the HBA manufacturer's instructions.

Clustering and Fabric zoning

If you plan to use clustering, install and configure the clustering software on the servers. Clustering is the organization of multiple servers into groups. Within a cluster, each server is a node. Multiple clusters compose a multi-cluster environment. The following example shows a multi-cluster environment with three clusters, each containing two nodes. The nodes share access to the disk array.

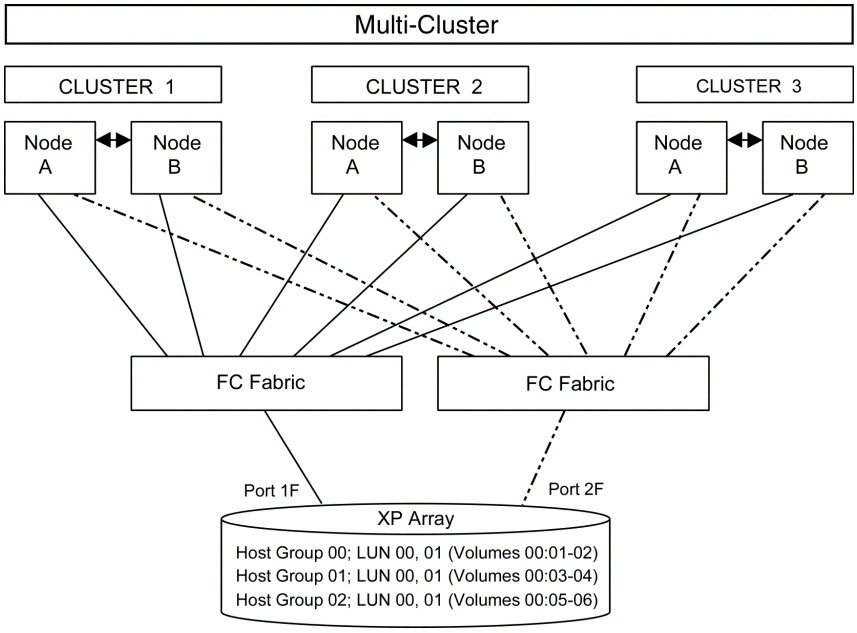

Within the Storage Area Network (SAN), the clusters may be homogeneous (all the same operating system) or they may be heterogeneous (mixed operating systems). How you configure LUN Security and fabric zoning depends on the operating system mix and the SAN configuration.

Fabric zoning and LUN security for multiple operating systems

By using appropriate zoning and LUN security, you can connect multiple clusters of various operating systems to the same switch and fabric:

- Host zones must contain only homogeneous operating systems.
- Storage port zones may overlap if more than one operating system needs to share an array port.
- Heterogeneous operating systems may share an XP array port if you use Secure Manager and set the appropriate host group and mode; all others must connect to a dedicated XP array port.
- Use Secure Manager for LUN isolation when multiple hosts connect through a shared array port. Secure Manager provides LUN security by allowing you to restrict which LUNs each host can access.

Environment	OS Mix	Fabric Zoning	LUN Security
Standalone SAN (non-clustered)	homogeneous (a single OS type present in the SAN)	Not required	Must be used when multiple hosts connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	
Clustered SAN	homogeneous (a single OS type present in the SAN)	Not required	Must be used when multiple cluster nodes connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	
Multi-Cluster SAN	homogeneous (a single OS type present in the SAN)	Not required	Must be used when multiple cluster nodes connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	

Connect the disk array

Connect the disk array to the host as follows:

1. The HP service representative verifies operational status of the disk array channel adapters, LDEVs, and paths.
2. The HP representative connects the Fibre Channel cables between the disk array and the fabric or host.
3. Verify the ready status of the disk array and peripherals.

Defining the paths

Use Command View (shown) or LUN Configuration Manager to create paths (LUNs) between hosts and volumes in the disk array, also called LUN mapping. LUN mapping includes these tasks:

- Configuring ports
- Setting LUN security
- Creating host groups by operating system and setting their host modes
- Assigning host bus adapter WWNs to host groups.
- Mapping volumes to host groups (by assigning LUNs).

For details, see the Command View or LUN Configuration Manager guide. HP recommends that you note LUNS and their ports, WWNs, nicknames, and LDEVs for later use in verifying host and device configuration.

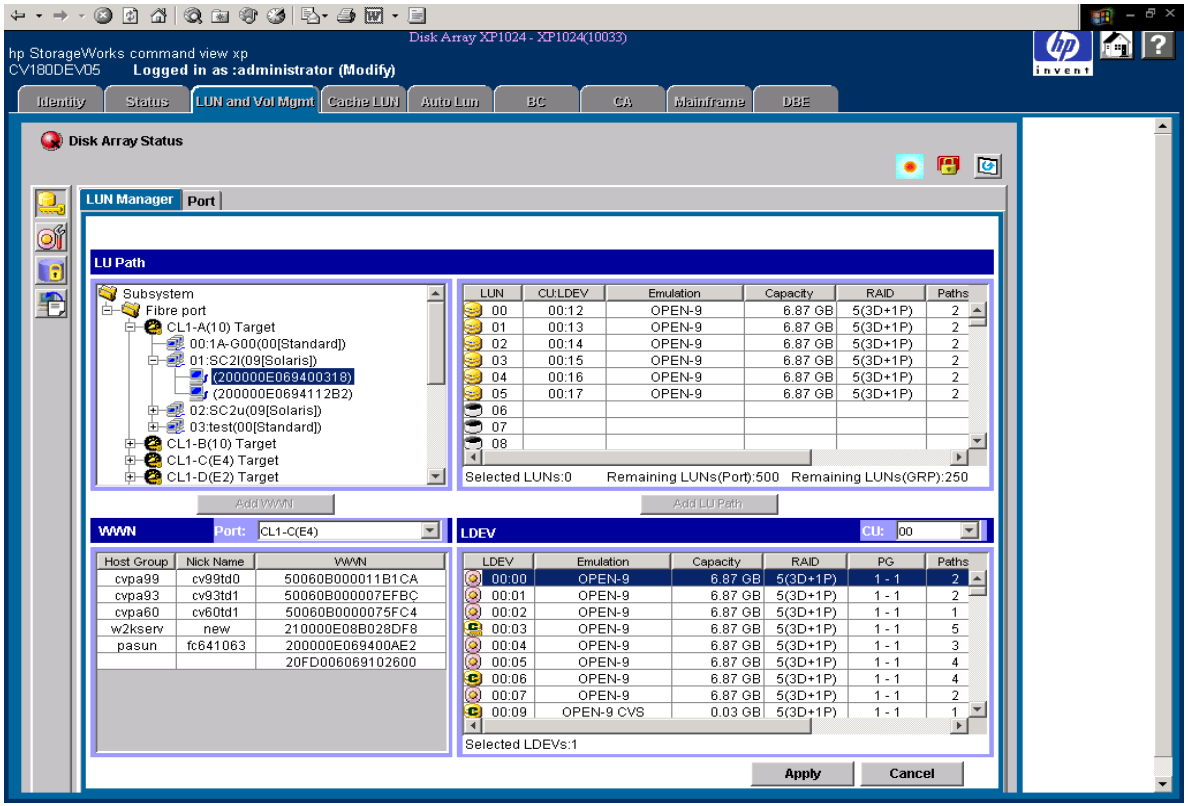

Verifying HBA installation

After configuring the ports on the disk array, verify that the HBAs are installed properly.

Use the **ioscan –f** command, and verify that the rows shown in the example are displayed. If these rows are not displayed, check the host adapter installation (hardware and driver installation) or the host configuration.

Example

```
# ioscan -f
Class    I  H/W Path         Driver    S/W State H/W Type   Description
=============================================================================================
...
fc       0  8/12             fcT1      CLAIMED   INTERFACE  HP Fibre Channel Mass Storage Adapter
lan      1  8/12.5           fcT1_cntl CLAIMED   INTERFACE  HP Fibre Channel Mass Storage Cntl
fcp      0  8/12.8           fcp       CLAIMED   INTERFACE  FCP Protocol Adapter
ext_bus  2  8/12.8.0.255.0   fcpdev    CLAIMED   INTERFACE  FCP Device Interface
...
```

Verifying device recognition

Verify that the HP-UX system recognizes the new devices on the disk array.

If the SCSI paths were defined after the system is powered on, you must halt and restart the system to allow the system to recognize the new devices.

To verify device recognition:

1. Log in to the system as **root**.
2. Display the device data to verify that the system recognizes the newly installed devices on the disk array. Use the **ioscan –fn** command to display the device data.

   On a system with a large LUN configuration, HP-UX may not build device files on all LUNs. Enter **insf –e** to build all missing device files.

Example

```
# ioscan -fn
class    I   H/W Path            Driver    S/W State H/W Type   Description
=============================================================================================
bc       6   14                  ccio      CLAIMED   BUS_NEXUS  I/O Adapter HP Fibre Channel Mass
fc       1   14/12               fcT1      CLAIMED   INTERFACE  Storage Adapter HP Fibre Channel Mass
lan      2   14/12.5             fcT1_cntl CLAIMED   INTERFACE  Storage Cntl /dev/fcms2
fcp      1   14/12.8             fcp       CLAIMED   INTERFACE  FCP Protocol Adapter HP A3308 FCP-SCSI MUX
ext_bus  6   14/12.8.0.0.0       fcpmux    CLAIMED   INTERFACE  Interface
target   9   14/12.8.0.0.0.0     tgt       CLAIMED   DEVICE
disk     4   14/12.8.0.0.0.0.0   sdisk     CLAIMED   DEVICE     HP OPEN-9 /dev/dsk/c6t0d0/dev/rdsk/c6t0d0
disk     5   14/12.8.0.0.0.0.1   sdisk     CLAIMED   DEVICE     HP OPEN-9*2 /dev/dsk/c6t0d1  /dev/rdsk/c6t0d1
ext_bus  7   14/12.8.0.255.0     fcpdev    CLAIMED   INTERFACE  FCP Device Interface
target   10  14/12.8.0.255.0.0   tgt       CLAIMED   DEVICE
ctl      5   14/12.8.0.255.0.0.0 sctl      CLAIMED   DEVICE     HP HPA3308 /dev/rscsi/c7t0d0
```

In the example:

HP OPEN-9 device: SCSI bus number = 14/12, bus instance = 6, SCSI target ID = 0, LUN = 0.

HP OPEN-9*2 device: SCSI bus number = 14/12, bus instance = 6, SCSI target ID = 0, LUN = 1.

If UNKNOWN is displayed for a disk, the HP 9000 system may not be configured properly. Refer to the HP documentation or contact HP customer support for assistance with the HP 9000 system or the HP-UX operating system.

3. Enter the device data for each disk array device in a table. See “Path worksheet” (page 50).
4. Construct the device file name for each device, using the device information, and enter the file names in your table. Use the following formula to construct the device file name:

   c*x***t***y***d***z***

   where
   x = SCSI bus instance number
   y = SCSI target ID
   z = LUN

   c stands for controller, **t** stands for target ID, and **d** stands for device. The numbers *x*, *y,* and *z* are hexadecimal.

Example

SCSI bus instance number	Hardware path	SCSI TID	LUN	File name
00	14/12.6.0	6	0	c6t0d0
00	14/12.6.1	6	1	c6t0d1

5. Verify that the SCSI TIDs correspond to the assigned port address for all connected ports (see mapping tables in appendix B, page 51, for values). If so, the logical devices are recognized properly.

   If the logical devices are not recognized properly:

- Check the AL-PA for each port using the LUN Manager software.
- If the same port address is set for multiple ports on the same loop (AL with HUB), all port addresses except one changed to another value, and the relationship between AL-PA and TID does not correspond to the mapping given in appendix B (page 51). Set a different address for each port, reboot the server, and then verify new device recognition again.
- If unused device information remains, the TID-to-AL-PA mapping will not correspond to the mapping given in appendix B (page 51). Renew the device information, and then verify new device recognition again.

Configure disk array devices

Configure the disk array devices in the same way you would configure any new disk on the host. Creating scripts to configure all devices at once may save you considerable time.

1. **For UNIX systems, configuring devices typically requires these steps:** Initialize, partition, and label each device using the disk management utility provided with the host server.
2. *(HP-UX only)* Create volumes (**pvcreate**), volume groups (**vgcreate)**, and logical volumes (**lvcreate)**.
3. Create a file system for each device using the **newfs** command.
4. Create a mount directory for each device using the **mkdir** command.
5. Enter each device into the mount table by editing **/etc/fstab**.
6. Reboot the host. After bootup use the **bdf** command to verify the devices auto-mounted.
7. Verify file system operation by copying a file to each device and then deleting it.

The HP-UX system uses the Logical Volume Manager (LVM) to manage the OPEN-*x* devices on the disk array. The instructions in this section do not explicitly cover all LVM configuration issues. For further information on LVM configuration, see the HP-UX user documentation.

HP System Administrator Manager (SAM) can be used instead of UNIX commands to configure SCSI disk devices. See appendix C (page 61) for further information.

Verifying the device files and drivers

The device files for new devices are created automatically during HP-UX startup. Each device must have a block-type device file in the **/dev/dsk** directory and a character-type device file in the **/dev/rdsk** directory. This procedure verifies both types of device files.

Some HP-compatible systems do not create the device files automatically. If verification shows that the device files were not created automatically, follow the instructions in "Creating the device files" (page 30) to create the device files manually.

To verify that the device files for the disk array devices were successfully created:

1. Display the block-type device files in the **/dev/dsk** directory using the **ls –l** command with the output piped to **more**. Verify there is one block-type device file for each disk array device.

Example

```
# ls –l /dev/dsk | more
Total 0
brw-r - - - - - 1 bin sys 28 0x006000 Dec 6 15:08 c6t0d0
brw-r - - - - - 1 bin sys 280 0x06100 Dec 6 15:08 c6t0d1
```

2. Verify that the block-type device file name for each device is correct.
3. Display the character-type device files in the **/dev/rdsk** directory using the **ls –l** command with the output piped to **more**. Verify that there is one character-type device file for each disk array device.

Example

```
# ls –l /dev/rdsk | more
Total 0
crw-r - - - - -  1 bin sys 177 0x006000 Dec 6 15:08 c6t0d0
crw-r - - - - -  1 bin sys 177 0x006100 Dec 6 15:08 c6t0d1
```

4. Use the device data table you created to verify that the character-type device file name for each device is correct.

   This task can also be accomplished with the **lssf** command.

5. After verifying the block-type and character-type device files, verify the HP-UX driver for the disk array using the **ioscan –fn** command.

Example

```
# ioscan -fn
Class    I  H/W Path            Driver    S/W State H/W Type   Description
=============================================================================================
bc       0                      root      CLAIMED   BUS_NEXUS
bc       1  8                   bc        CLAIMED   BUS_NEXUS  Bus Converter
fc       0  8/12                fcT1      CLAIMED   INTERFACE  HP Fibre Channel Mass Storage
fcp      0  8/12.8              fcp       CLAIMED   INTERFACE  FCP Protocol Adapter
ext_bus  2  8/12.8.0.255.0      fcpdev    CLAIMED   INTERFACE  FCP Device Interface
disk     3  8/12.8.8.255.0.6.0  sdisk     CLAIMED   DEVICE     HITACHI OPEN-9
                    /dev/dsk/c2t6d0      /dev/rdsk/c2t6d0
disk     4  8/12.8.8.255.0.6.1  sdisk     CLAIMED   DEVICE     HITACHI OPEN-9
                    /dev/dsk/c2t6d1      /dev/rdsk/c2t6d1
disk     5  8/12.8.8.255.0.8.0  sdisk     CLAIMED   DEVICE     HITACHI 3390*3B
                    /dev/dsk/c2t8d0      /dev/rdsk/c2t8d0
:
#
```

Creating the device files

If the device files were not created automatically when the system was restarted, issue the **insf –e** command in the **/dev** directory to create the device files. After this command is executed, repeat the procedures in the previous section to verify new device recognition and the device files and driver.

Example

```
# insf -e
insf: Installing special files for mux2 instance 0 address 8/0/0
       :          :          :         :
       :          :          :         :
#
```

Failure of the **insf –e** command indicates a SAN problem.

If the device files for the new disk array devices cannot be created automatically, you must create the device files manually using the **mknod** command as follows:

1. Retrieve the device information you recorded earlier.

2. Construct the device file name for each device, using the device information, and enter the file names in your table. Use the following formula to construct the device file name:

   c*x***t***y***d***z***

   where
   x = SCSI bus instance number
   y = SCSI target ID
   z = LUN

   c stands for controller, **t** stands for target ID, and **d** stands for device. The numbers *x*, *y,* and *z* are hexadecimal.

3. Construct the minor number for each device, using the device information, and enter the file names in your table. Use the following formula to construct the minor number:

   0x*xxyz*00

   where

   xx = SCSI bus instance number
   y = SCSI target ID
   z = LUN

4. Display the driver information for the system using the **lsdev** command.

Example

```
# lsdev
Character   Block   Driver   Class
    :         :       :        :
   188        31     sdisk    disk
#
```

5. Enter the major numbers for the device drivers into the table. You should now have all required device and driver information in the table.

6. Create the device files for all disk array devices (SCSI disk and multiplatform devices) using the **mknod** command. Create the block-type device files in the **/dev/dsk** directory and the character-type device files in the **/dev/rdsk** directory.

Example

```
# cd /dev/dsk Go to /dev/dsk directory.
# mknod /dev/dsk/c2t6d0 b 31 0x026000
Create block-type file.
```

```
File name, b=block-type, 31=major #, 0x026000= minor #
# cd /dev/rdskGo to /dev/rdsk directory.
# mknod /dev/rdsk/c2t6d0 c 188 0x026000
Create character-type file.
File name, c=character-type, 177=major #, 0x026000=minor #
 :
#
```

The character-type device file is required for volumes used as raw devices (for example, 3390-3A/B/C). The block-type device file is not required for volumes used as raw devices.

If you need to delete a device file, use the **rm –i** command.

Example of a completed device data table

Bus no.	Instance (XX)	Disk no.	H/W path	Driver	Device type	TID (Y)	LUN (Z)	Device file	Minor # 0xXXYZ00	Major # for char. files	Major # for block files
8/12	02	3	8/12.8.8.255.0.6.0	sdisk	OPEN-9	6	0	c2t6d0	0x026000	188	31
8/12	02	4	8/12.8.8.255.0.6.1	sdisk	OPEN-9	6	1	c2t6d1	0x026100	188	31
8/12	02	5	8/12.8.8.255.0.8.0	sdisk	3390-3 B	8	0	c2t8d0	0x028000	188	31

Creating the physical volumes

A physical volume must be created for each new SCSI disk device.

To create the physical volumes:

1. Use the **pvcreate** command to create the physical volumes with the character-type device file as the argument. Specify the **/dev/rdsk** directory.

Example

```
# pvcreate /dev/rdsk/c6t0d0
Physical volume "/dev/rdsk/c6t0d0" has been successfully created.
  :
# pvcreate /dev/rdsk/c6t0d1
```

```
Physical volume "/dev/rdsk/c6t0d1" has been successfully created.
```

Do not use the **–f** option with the **pvcreate** command. This option creates a new physical volume forcibly and overwrites the existing volume. If you accidentally enter the character-type device file for an existing volume, you will lose the data on that volume.

2. Repeat step 1 for each OPEN-*x* device on the disk array.

Creating new volume groups

You must create new volume groups for the new physical volumes. If desired, you can also add any of the volumes on the disk array to existing volume groups using the **vgextend** command. The physical volumes that make up one volume group can be located either in the same disk array or in other disk arrays.

To allow more volume groups to be created, use SAM to modify the HP-UX system kernel configuration. See appendix C (page 61) for details.

To create volume groups:

1. Use the **vgdisplay** command to display the existing volume groups.
2. Choose a unique name for the new volume group (for example: vg06).
3. Create the directory for the new volume group.

Example

```
# mkdir /dev/vg06
```

4. Use the **ls –l** command (with the output piped to **grep** to display only the files containing “group”) to display the minor numbers for the existing group files.

Example

```
# ls -1 /dev/vg* | grep group
crw-rw-rw 1 root root 64  0x0000000  Nov7 08:13 group
  :
```

5. Choose a unique minor number for the new group file in sequential order (for example, when existing volume groups are vg00-vg05 and the next group name is vg06, use minor number 06 for the vg06 group file).

   The minor numbers are hexadecimal (for example, the 10th minor number is 0x0a0000).

6. Use **mknod** to create the group file for the new directory. Specify the volume group name, major number, and minor number. The major number for all group files is 64.

Example

In this example: group name = vg06, major number of group file = 64, minor number of existing group file = 06 (which must be unique for each volume group), and c = character.

```
# mknod /dev/vg06/group c 64 0x060000
  :
```

7. Create the volume group.

   To allocate more than one physical volume to the new volume group, add the other physical volumes, separated by a space.

Example

```
# vgcreate /dev/vg06 /dev/dsk/c6t0d0
Volume group "/dev/vg06" has been successfully created.
Volume group configuration for /dev/vg06 has been saved in
/etc/1vmconf/vg06.conf.
```

   For Logical Unit Size Expansion (LUSE) volumes that contain more than 17 OPEN-8/9 LDEVs or more than 7043 MB (OPEN-8/9*n-CVS), use the **–s** and **–e** physical extent (PE) parameters of the **vgcreate** command. See “LUSE device parameters” (page 56).

   If you need to delete a volume group, use the **vgremove** command (for example, **vgremove /dev/vgnn**). If the **vgremove** command does not work because the volume group is not active, use the **vgexport** command (for example, **vgexport /dev/vgnn**).

8. Use the **vgdisplay** command to verify that the new directory was created.

9. Use **vgdisplay –v** to verify that the volume group was created correctly. The **–v** option displays the detailed volume group information.

Example

```
# vgdisplay –v /dev/vg06
- - -  Volume groups  - - -
VG Name/dev/vg06
VG Write Accessread/write
VG Statusavailable
Max LV255
Cur LV0
Open LV0
Max PV16
```

```
Cur PV1
Act PV1
Max PE per PV1016
VGDA2
PE Size (Mbytes)4
Total PE586
Alloc PE0
Free PE586
Total PVG0

- -  Physical Volumes  - - -
PV Name/dev/dsk/c6t0d0
PV Statusavailable
Total PE586
Free PE586
```

Creating logical volumes

For logical volume configuration, use these commands:

lvremove Deletes a logical volume.
Any file system attached to the logical volume must be unmounted before executing the **lvremove** command.

Example: **`lvremove /dev/vgnn/lvolx`**

lvextend Increases the size of an existing logical volume.

Example: **`lvextend –L size /dev/vgnn/lvolx`**

lvreduce Decreases the size of an existing logical volume.

Any file system attached to the logical volume must be unmounted before executing the **lvreduce** command.

Example: **`lvreduce –L size /dev/vgnn/lvolx`**

Caution *Data within the file system can be lost after execution of **lvreduce**.*

Create logical volumes after you create volume groups. A logical volume must be created for each new SCSI disk device

To create logical volumes:

1. Use the **lvcreate –L** command to create a logical volume.

   Specify the volume size (in megabytes) and the volume group for the new logical volume. HP-UX assigns the logical volume numbers automatically (lvol1, lvol2, lvol3). Use the following capacity values for the size parameter:

   OPEN-K = 1740
   OPEN-3 = 2344
   OPEN-8 = 7004
   OPEN-9 = 7004
   OPEN-E = 13888
   OPEN-L = 34756
   OPEN-V = 61432

   To calculate S1 for CVS, LUSE, and CVS LUSE volumes, first use the **vgdisplay** command to display the physical extent size (PE Size) and usable number of physical extents (Free PE) for the volume. Calculate the maximum size value (in MB) as follows:

   S1 = (PE Size) × (Free PE)

   Logical volumes can span multiple physical volumes. Use the **diskinfo** command for extended LUNs.

Example

Create an OPEN-3 logical volume the size of the physical volume, using 2344 for the size parameter. An OPEN-9 volume uses 7040 for the size parameter to create a logical volume the size of the physical volume.

```
# lvcreate –L 2344 /dev/vg06
Logical volume "/dev/vg06/lvol1" has been successfully created
with character device "/dev/vg06/rlvol1".
Logical volume "/dev/vg06/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg06 has been saved in
/etc/1vmconf/vg06.cof.
```

2. Use the **lvdisplay** command to verify that the logical volume was created correctly.

Example

```
# lvdisplay /dev/vg06/lvol1
- - -  Logical volume - - -
LV Name                /dev/vg06/lvol1
VG Name                /dev/vg06
LV Permission          read/write
LV Status              available/syncd
Mirror copies          0
Consistency Recovery   MWC
Schedule               parallel
LV Size (Mbytes)       2344
Current LE             586
Allocated PE           586
Stripes                0
Stripe Size (Kbytes)   0
Bad block              on
Allocation             strict
```

3. Repeat steps 1 and 2 for each logical volume to be created.

   You can create only one logical volume at a time. However, you can verify multiple logical volumes at a time.

Creating the file systems

Create the file system for each new logical volume on the disk array. The default file system types are:

- HP-UX OS version 10.20 = **hfs** or **vxfs**, depending on entry in the **/etc/defaults/fs** file.
- HP-UX OS version 11.0 = **vxfs**
- HP-UX OS version 11.i = **vxfs**

To create file systems:

1. Use the **newfs** command to create the file system using the logical volume as the argument.

Example

```
# newfs /dev/vg06/rlvol1
newfs: /etc/default/fs is used for determining the file system
type
mkfs  (hfs): Warning -272 sectors in the last cylinder are not
allocated.
```

```
mkfs  (hfs): /dev/vg06/rlvol1 - 2400256 sectors in 3847 cylinders
of 16 tracks, 2547.9MB in 241 cyl groups (16 c/g, 10.22Mb/g,
1600 ig/g)
Super block backups (for fsck -b) at:
16,   10040,  20064,  30038,  40112,  50136,  60160,  70184,
80208,  90232, . . . 2396176
```

Example

```
# newfs /dev/vg06/rlvol1Create file system
newfs: / etc/default/fs is used for determining the file system
type
mkfs (hfs): ...
:
7188496, 7198520, 7208544
#
```

Example

```
# newfs -F vxfs /dev/vg06/rlvol1Specify file system type
:
# newfs -F hfs /dev/vg06/rlvol2
```

2. Repeat step 1 for each logical volume on the disk array.

Setting the I/O timeout parameter

The I/O timeout value for each disk device on the system must be set to 60 seconds.

To change the I/O time-out parameter:

1. Verify the current I/O time-out value using the **pvdisplay** command:

Example

```
# pvdisplay /dev/dsk/c0t6d0
This causes current values to be displayed:
--- Physical volumes ---
PV Name                /dev/dsk/c0t6d0
VG Name                /dev/vg06
PV Status              available
Allocatable            yes
VGDA                   2
Cur LV                 1
PE Size (Mbytes)       4
Total PE               586
Free PE                0
Allocated PE           586        [OPEN-9]
Stale PE               0
IO Timeout (Seconds)   default    [I/O timeout value]
```

2. If the I/O timeout value is not 60, change the value to 60 using the **pvchange -t** command:

Example

```
# pvchange -t 60 /dev/dsk/c0t6d0
```

This result is displayed:

```
Physical volume "/dev/dsk/c0t6d0" has been successfully changed.
Volume Group configuration for /dev/vg06 has been saved in
/etc/lvmconf/vg06.cof
```

3. Verify that the new I/O timeout value is 60 seconds using the **pvdisplay** command:

Example

```
# pvdisplay /dev/dsk/c0t6d0
```

Current values are now displayed, confirming that the new I/O timeout value is 60 seconds:

```
--- Physical volumes ---
PV Name               /dev/dsk/c0t6d0
VG Name               /dev/vg06
PV Status             available
 :
Stale PE              0
IO Timeout (Seconds)  60    [New I/O timeout value]
```

4. Repeat the steps above for each new disk that is connected to the system.

Creating the mount directories

Create a mount directory for each logical volume. Choose a unique name for each mount directory that identifies the logical volume.

Example The following example shows the root directory as the location for the mount directories.

To create a mount directory for each logical volume:

1. Use **mkdir** with the new mount directory name as the argument to create the mount directory.

   ```
   # mkdir /AHPMD-LU00
   ```

2. Use the **ls –x** command to verify the new mount directory.

   ```
   # ls -x
   AHPMD-LU00   bin    dev        device   etc      export
   floppy       home   hstsboof   kadb     kernel   lib
   ```

3. Repeat steps 1 and 2 for each logical volume on the disk array.

Mounting and verifying the file systems

After the mount directories have been created, mount and verify the file system for each logical volume.

To mount and verify the file systems:

1. Use **mount** to mount the file system for the volume.

Example

```
# mount /dev/vg06/lvol1  /AHPMD-LU00
```

Repeat this step for each logical volume on the disk array.

If you need to unmount a file system, use the **unmount** command.

2. Use the **bdf** command to verify that the file systems are correct. The capacity is listed under Kbytes.

Example

```
# bdf
Filesystem          Kbytes    used   avail %used Mounted on
/ldev/vg00/lvol1     59797   59364       0  100%    /
   :
/ldev/vg06/lvol1   2348177       9 2113350    0%
/AHPMD-LU00
/ldev/vg07/lvol1   2348177       9 2113350    0%
/AHPMD-LU01
/ldev/vg08/lvol1   7052764       9 6347478    0%
/AHPMD-LU02
```

3. As a final verification, perform some basic UNIX operations (for example file creation, copying, and deletion) on each logical device to make sure that the devices on the disk array are fully operational.

Example

```
#cd /AHPMD-LU00
#cp /bin/vi /AHPMD-LU00/vi.back1
#ls -l
drwxr-xr-t    2  root   root      8192 Mar  15  11:35  lost+found
-rwxr-xr-x    1  root   sys     217088 Mar  15  11:41  vi.back1
#cp vi.back1 vi.back2
#ls -l
drwxr-xr-t    2  root   root      8192 Mar  15  11:35  lost+found
-rwxr-xr-x    1  root   sys     217088 Mar  15  11:41  vi.back1
-rwxr-xr-x    1  root   sys     217088 Mar  15  11:52  vi.back2
```

Setting and verifying the auto-mount parameters

Set up and verify the auto-mount parameters for each new volume. The **/etc/checklist** file (which can also be called the **/etc/fstab** file) contains the auto-mount parameters for the logical volumes.

To set up and verify the auto-mount parameters:

1. Edit the **/etc/checklist** (**/etc/fstab**) file to add a line for each OPEN-*x* device on the disk array. This example and the following table show the auto-mount parameters.

Example

```
#cp -ip /etc/checklist  /etc/checklist.standard
#vi /etc/checklist
/dev/vg00/lvol1     /              hfs      rw          0      1       #  root
/dev/vg00/lvol2     swap           ignore   rw          0      0       #  primary swap
:
/dev/vg06/lvol1     /AHPMD-LU00    hfs      defaults    0      2       #  AHPMD-LU00
/dev/vg06/lvol2     /AHPMD-LU01    hfs      defaults    0      2       #  AHPMD-LU01
P1                  P2             P3       P4          P5     P6      P7
```

Parameter	Name	Enter
P1	Device to mount	Block-type device file name
P2	Mount point	Mount directory name
P3	File system	Type of file system (for example, hfs, vxfs)
P4	Mount options	“defaults” or other appropriate mount options
P5	Enhance	0
P6	File system check (fsck pass)	Order for performing file system checks
P7	Comments	Comment statement

2. Reboot the system.
3. Use the **bdf** command to verify the file system again.

2

Troubleshooting

If you encounter an error condition, see “Error conditions” on page 44 for recommended actions.

If you are unable to resolve an error condition, ask your HP support representative for assistance. See “Calling the HP support center” on page 47.

Error conditions

Error Condition	Recommended Action
The logical devices are not recognized by the host.	Verify that the READY indicator lights on the disk array are ON. Verify that fiber cables are correctly installed and firmly connected. Verify that the target IDs are properly configured. The LUNs for each TID must start at 0 and continue sequentially without skipping any numbers. Verify that the TIDs/WWNs on each bus are unique. Do not install two devices with the same ID on the same bus. Recheck the buses for new devices. Verify that LUSE devices are not intermixed with normal LUNs on the same port. Verify that the maximum number of LUSE devices per port has not been exceeded. Verify that the disk array Host Mode is set correctly.
The host does not reboot properly after hard shutdown.	If you power off the host without executing the shutdown process, wait three minutes to allow the disk array's internal timeout process to purge queued commands. If the host restarts while the disk array is processing queued commands, the host may not reboot successfully.
Physical volumes cannot be created (**pvcreate**).	Verify that the disk array logical devices are correctly formatted. Verify that the character-type device file exists. Verify that the correct character-type device file name is used with **pvcreate** (for example, **/dev/rdsk/...**).

Error Condition	Recommended Action
Volume group cannot be created (**vgcreate**).	Verify that the directory for the new volume group exists. Verify that the control file exists. Verify that the correct major number (64) and minor number are used with **mknod**. Verify that the block-type device file exists and is entered correctly with **vgcreate**. Verify that the physical volume is not already allocated to another volume group.
Logical volumes cannot be created (**lvcreate**).	Verify that the volume capacity for OPEN-x volumes is not greater than the maximum capacity allowed. See the Device Emulations Appendix. Verify that the capacity of the volume group is not less than the total capacity of the partitioned logical volume.
The file system cannot be created (**newfs** command).	Verify that the character-type device file is entered correctly with **newfs** (for example, **/dev/vg01/r/vo/x**).
A file system is not mounted after rebooting.	Verify that the host was restarted correctly. Verify that the file system attributes are correct. Verify that the auto-mount information in the /**etc/checklist** (**/etc/fstab**) file is correct.
The disk array performs a self reboot because the disk array was busy or it logged a panic message.	Reboot the host.
The disk array responds “Not Ready” or the disk array has displayed “Not Ready” and timed out.	Contact HP.

Error Condition	Recommended Action
The host detects a parity error.	Check the HBA and make sure it was installed properly. Reboot the host.
The host hangs or devices are declared and the host hangs.	Make sure there are no duplicate disk array TIDs and that disk array TIDs do not conflict with any host TIDs.

Calling the HP support center

If you need to call HP customer support, provide as much information about the problem as possible, including the circumstances of the error or failure and the exact content of any error messages.

Depending on your system configuration, you may be able to view error messages as follows:

- View SIMs in Command View (Device Health tab).
- View R-SIMs in Remote Control XP, including reference codes and severity levels of recent R-SIMs.
- View SIMs that generate SNMP traps on the host.

A

Worksheets

Path worksheet

Bus no.	Instance (XX)	Disk no.	H/W path	Driver	Device type	TID (Y)	LUN (Z)	Device file	Minor # 0xXXYZ00	Major # for char. files	Major # for block files

B

Disk array device emulations

This appendix provides information about disk array supported emulations and device type specifications. Some parameters may not be relevant to your array. Consult your HP representative for information about supported configurations for your system.

Supported emulations

XP Type	Emulation	OPEN-x	LUSE	CVS	LUSE & CVS
XP48 XP512	OPEN-3	Yes	Yes	Yes	Yes
	OPEN-8	Yes	Yes	Yes	Yes
	OPEN-9	Yes	Yes	Yes	Yes
	OPEN-E	Yes	Yes	Yes	Yes
	OPEN-K	Yes	Yes	Yes	Yes
	OPEN-L	Yes	Yes		
	OPEN-M	Yes	Yes		
	OPEN-V				
XP128 XP1024 XP12000	OPEN-3	Yes	Yes	Yes	Yes
	OPEN-8	Yes	Yes	Yes	Yes
	OPEN-9	Yes	Yes	Yes	Yes
	OPEN-E	Yes	Yes	Yes	Yes
	OPEN-K				
	OPEN-L	Yes	Yes		
	OPEN-M				
	OPEN-V	Yes	Yes		

Device type specifications

Device Type *(Note 1)*	Category *(Note 2)*	Blocks (512 bytes)	Sector Size (bytes)	# of Cylinders	Heads	Sectors per Track	Capacity MB* *(Note 3)*
OPEN-3	SCSI disk	4806720	512	3338	15	96	2347
OPEN-8	SCSI disk	14351040	512	9966	15	96	7007
OPEN-9	SCSI disk	14423040	512	10016	15	96	7042
OPEN-E	SCSI disk	28452960	512	19759	15	96	13893
OPEN-L	SCSI disk	71192160	512	49439	15	96	34761
OPEN-V	SCSI disk	max=125827200	512	Note 5	15	128	Note 6
LUSE							
OPEN-3*n	SCSI disk	4806720*n	512	3338*n	15	96	2347*n
OPEN-8*n	SCSI disk	14351040*n	512	9966*n	15	96	7007*n
OPEN-9*n	SCSI disk	14423040*n	512	10016*n	15	96	7042*n
OPEN-E*n	SCSI disk	28452960*n	512	19759*n	15	96	13893*n
OPEN-L*n	SCSI disk	71192160*n	512	49439*n	15	96	34761*n
OPEN-V*n	SCSI disk	max=125827200 Note 4	512	Note 5	15	128	Note 6
CVS							
OPEN-3 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-8 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-9 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-E CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
CVS LUSE							
OPEN-3*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-8*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-9*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-E*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-V*n	SCSI disk	Note 4	512	Note 5	15	128	Note 6

*Capacity = (512 x number of blocks) ÷ 1024^2

Note 1: The availability of a disk type depends on the disk array.

Note 2: The devices are defined to the host as SCSI disk devices, even though the interface is Fibre Channel.

Note 3 The device capacity can sometimes be changed by the BIOS or host adapter board. This may make actual capacity different from that listed in the table.

Note 4: The number of blocks for a CVS volume is calculated as follows:
# of blocks = (# of cylinders) × (# of heads) × (# of sectors per track)

Example 1: For an OPEN-3 CVS volume with capacity = 37 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (96 sectors per track) = 76320

Example 2: For an OPEN-V CVS volume with capacity = 49 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (128 sectors per track) = 101760

Note 5: The number of cylinders for a CVS volume is calculated as follows (↑ …↑ means that the value should be rounded up to the next integer):

OPEN-3/8/9/E: The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑

Example: For an OPEN-3 CVS volume with capacity = 37 MB:
# of cylinders = ↑ 37 × 1024/720 ↑ = ↑ 52.62 ↑ (rounded up to next integer) = 53 cylinders

OPEN-V: The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

Example: For an OPEN-V CVS volume with capacity = 49 MB:
# of cylinders = ↑ 49 × 16/15 ↑ = ↑ 52.26 ↑ (rounded up to next integer) = 53 cylinders

OPEN-3/8/9/E: The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑ × n

Example: For a CVS LUSE volume with capacity = 37 MB and n = 4
# of cylinders = ↑ 37 × 1024/720↑ × 4 = ↑ 52.62↑ × 4 = 53 × 4 = 212

OPEN-V: The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑ × n

Example: For an OPEN-V CVS LUSE volume with capacity = 49 MB and n = 4
# of cylinders = ↑ 49 × 16/15↑ × 4 = ↑ 52.26↑ × 4 = 53 × 4 = 212

Note 6: The capacity of an OPEN-3/8/9/E CVS volume is specified in MB, not number of cylinders. The capacity of an OPEN-V CVS volume can be specified in MB or number of cylinders. You set the volume size using the LUN Configuration Manager or Command View software.

LUSE device parameters

Device type		Physical extent size (PE)	Max physical extent size (MPE)
OPEN-K/3/8/9/E OPEN-3/K*n (n= 2 to 36) OPEN-3/K-CVS OPEN-3/K*n-CVS (n = 2 to 36)		default	default
OPEN-8/9*n	n = 2 to 17	default	default
	n = 18	8	15845
	n = 19	8	16725
	n = 20	8	17606
	n = 21	8	18486
	n = 22	8	19366
	n = 23	8	20247
	n = 24	8	21127
	n = 25	8	22007
	n = 26	8	22888
	n = 27	8	23768
	n = 28	8	24648
	n = 29	8	25529
	n = 30	8	26409
	n = 31	8	27289
	n = 32	8	28170
	n = 33	8	29050
	n = 34	8	29930
	n = 35	8	30810
	n = 36	8	31691
OPEN-E*n	n = 2 to 9	default	default
	n = 10	8	17366

Device type		Physical extent size (PE)	Max physical extent size (MPE)
	n = 11	8	19102
	n = 12	8	20839
	n = 13	8	22576
	n = 14	8	24312
	n = 15	8	26049
	n = 16	8	27786
	n = 17	8	29522
	n = 18	8	31259
	n = 19	8	32995
	n = 20	8	34732
	n = 21	8	36469
	n = 22	8	38205
	n = 23	8	39942
	n = 24	8	41679
	n = 25	8	43415
	n = 26	8	45152
	n = 27	8	46889
	n = 28	8	48625
	n = 29	8	50362
	n = 30	8	52098
	n = 31	8	53835
	n = 32	8	55572
	n = 33	8	57308
	n = 34	8	59045
	n = 35	8	60782
	n = 36	8	62518
OPEN-L*n	n = 2 to 3	default	default

Device type		Physical extent size (PE)	Max physical extent size (MPE)
OPEN-8/9/E-CVS **OPEN-V**		default	default
OPEN-8/9/E*n-CVS **OPEN-V*n** **(n = 2 to 36)**	70 to 119731(MB) × N1	8	default
	119732 to (MB) × N1	8	N2

N1 = [VCS volume capacity (in MB) from Remote Console PC] × n

N2 = ↑ N1 / PE ↑ (↑ means round up to next integer)

Example: CVS volume capacity is 6000 MB for OPEN-9*22-CVS volume:

N1 = 6000 × 22 = 132000

N2 = ↑ 132000/8 ↑ = 16500

SCSI TID map for Fibre Channel adapters

When an arbitrated loop (AL) is established or reestablished, the port addresses are assigned automatically to prevent duplicate TIDs. With the SCSI over Fibre Channel protocol (FCP), there is no longer a need for target IDs in the traditional sense.

SCSI is a bus-oriented protocol requiring each device to have a unique address since all commands go to all devices. For Fibre Channel, the AL-PA is used instead of the TID to direct packets to the desired destination.

Unlike traditional SCSI, when control of the loop is acquired, a point-to-point connection is established from initiator to target. To enable transparent use of FCP, the operating system maps a TID to each AL-PA.

The host maps SCSI protocol to Fibre Channel protocol and detects and accesses Fibre Channel-connected devices using device files (**/dev/dsk/c*t*d*** and **/dev/rdsk/c*t*d***) in the same way as for SCSI-connected devices. The device files for Fibre Channel-connected devices are configured in a different way from SCSI-connected devices, because Fibre Channel supports 126 addresses per path while SCSI supports 16 TIDs per path.

Table 1 (page 60) identifies the fixed mappings between the TID (drive) values assigned by the operating system and the Fibre Channel native addresses (AL_PA/SEL_ID) for Fibre Channel adapters. The controller number (the **dks** value in **/dev/dsk/dks*d*l*s***) depends on the server configuration, and a different value is assigned per each column of table 1 (page 60).

The mapping cannot be guaranteed under the following conditions:

- when disk array devices and other types of devices are connected in the same loop
- when information for unused devices remains in the server system
- when multiple ports participate in the same arbitrated loop

Table 1. AL-PA to SCSI TID mapping (t value)

AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value
EF	0	CD	0	B2	0	98	0	72	0	55	0	3A	0	25	0
E8	1	CC	1	B1	1	97	1	71	1	54	1	39	1	23	1
E4	2	CB	2	AE	2	90	2	6E	2	53	2	36	2	1F	2
E2	3	CA	3	AD	3	8F	3	6D	3	52	3	35	3	1E	3
E1	4	C9	4	AC	4	88	4	6C	4	51	4	34	4	1D	4
E0	5	C7	5	AB	5	84	5	6B	4	4E	5	33	5	1B	5
DC	6	C6	6	AA	6	82	6	6A	6	4D	6	32	6	18	6
DA	7	C5	7	A9	7	81	7	69	7	4C	7	31	7	17	7
D9	8	C3	8	A7	8	80	8	67	8	4B	8	2E	8	10	8
D6	9	BC	9	A6	9	7C	9	66	9	4A	9	2D	9	0F	9
D5	10	BA	10	A5	10	7A	10	65	10	49	10	2C	10	08	10
D4	11	B9	11	A3	11	79	11	63	11	47	11	2B	11	04	11
D3	12	B6	12	9F	12	76	12	5C	12	46	12	2A	12	02	12
D2	13	B5	13	9E	13	75	13	5A	13	45	13	29	13	01	13
D1	14	B4	14	9D	14	74	14	59	14	43	14	27	14	00	–
CE	15	B3	15	9B	15	73	15	56	15	3C	15	26	15		

C

Reference information for SAM

The HP System Administrator Manager (SAM) allows you to perform HP-UX system administration functions, including:

- setting up users and groups
- configuring the disks and file systems
- performing auditing and security activities
- editing the system kernel configuration

This appendix provides instructions for:

- using SAM to configure the disk devices
- using SAM to set the maximum number of volume groups

Configuring the devices using SAM

The SAM Areas window displays the system administration functions and allows you to select the desired function. The **Disks and File Systems** function allows you to configure new disk devices for LVM operations.

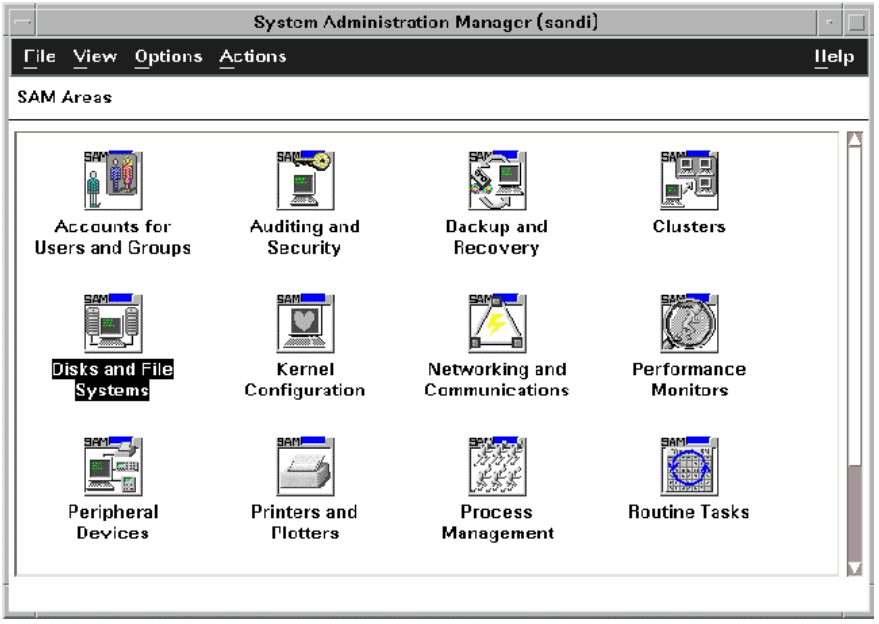

To configure the newly installed SCSI disk devices:

1. Select **Disks and File Systems**, then select **Disk Devices**.

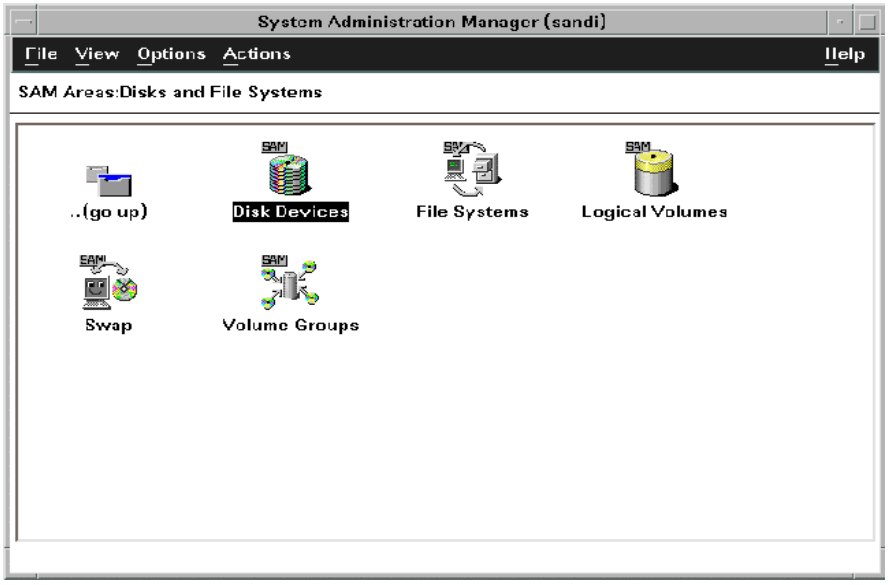

2. Verify that the new disk array devices are displayed in the **Disk Devices** window.
3. Select the device to configure, select the **Actions** menu, select **Add**, and then select **Using the Logical Volume Manager**.
4. In the **Add a Disk Using LVM** window, select **Create...** or **Extend a Volume Group**.
5. In the **Create a Volume Group** window, enter the name of the new or existing volume group to assign the new device to, and then click **OK** twice. The **Add a Disk Using LVM** window now displays the volume group name.
6. Select **Add New Logical Volume** to open the **Create New Logical Volumes** window.
7. In the **Create New Logical Volumes** window, enter the name, size in megabytes, and mount directory for the new logical volume. Select **Add**, and then click **OK** twice.
8. Repeat steps 3 through 7 for each new disk device on the disk array.

Setting the maximum number of volume groups using SAM

The HP-UX kernel specified the maximum number of volume groups that can be created. The default is 10. You may need to allows you to change this number to accommodate new devices on the disk array.To change this number, use the Kernel Configuration function.

To change the maximum number of volume groups:

1. Select the **Kernel Configuration** function, then select **Configurable Parameters**.
2. In the **Configurable Parameters** window, select the **maxvgs** parameter, select the **Actions** menu, and then select **Modify Configurable Parameter**.
3. In the **Modify Configurable Parameter** window, enter the desired maximum number of volume groups in the **Formula/Value** field, and then click **OK**.
4. In the **Configurable Parameters** window, make sure that none of the parameters is selected. Then select the **Actions** menu, and select **Create New Kernel**.
5. When the configuration window opens, click **Yes** to create the new kernel (or click **No** to return to the **Configurable Parameters** window.)
6. When the **Reboot the System** window opens, click **OK** to move the new kernel into place and reboot the system.

Glossary

AL	Arbitrated loop.
AL-PA	Arbitrated loop physical address.
BC	HP StorageWorks Business Copy XP. BC lets you maintain up to nine local copies of logical volumes on the disk array.
CA	HP StorageWorks Continuous Access XP. CA lets you create and maintain duplicate copies of local logical volumes on a remote disk array.
Command View	HP StorageWorks Command View XP, a software product for managing XP arrays. Command View runs on a Windows-based management workstation.
command device	An LDEV that transfers RAID commands to BC or CA logical volumes.
CVS	CVS devices (OPEN-x CVS) are custom volumes that are smaller than normal fixed-sized logical disk devices (volumes).
DKC (disk controller unit)	The array cabinet that houses the channel adapters and service processor (SVP).
DKU (disk cabinet unit)	The array cabinets that house the disk array physical disks.
emulation modes	Emulation modes can be assigned to LDEVs to make them operate like standard OPEN system disk drives. The emulation mode of an LDEV determines its capacity. Refer to the appendices for device capacities.
FC	Fibre Channel.

FC-AL	Fibre Channel arbitrated loop.
FCP	Fibre Channel Protocol.
HBA	Host bus adapter.
HP	Hewlett-Packard Company.
LDEV	Logical device. An LDEV is created when a RAID group is divided into sections using a selected host emulation mode (for example, OPEN-9 or OPEN-M). The number of resulting LDEVs depends on the emulation mode. "LDEV" and "volume" are synonyms.
LUN	Logical unit number. A LUN results from mapping a SCSI logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN. For example, a LUN associated with two OPEN-3 LDEVs has a size of 4,693 MB.
LUSE	Logical Unit Size Expansion, a feature which logically combines LDEVs so they appear as a larger LDEV. This allows a LUN to be associated with 2 to 36 LDEVs. LUSE allows applications to access data requiring large amounts of disk space.
OFC	Open Fibre Control.
OPEN-*x*	A general term describing any one of the supported OPEN emulation modes (for example, OPEN-L).
OS	Operating system.
PA	Physical address.
path	"Path" and "LUN" are synonymous. Paths are created by associating a port, a target, and a LUN ID with one or more LDEVs.
port	A connector on a channel adapter card in the disk array. A port passes data between the disk array and external devices, such as a host server. Ports are named using a port group and port letter, for example, CL1-A.

RAID	Redundant array of independent disks.
remote console PC	The PC running HP StorageWorks Remote Control XP.
Remote Control (RC)	HP StorageWorks Remote Control XP. A software product used for managing XP arrays.
R-SIM	Remote service information message.
SCSI	Small computer system interface.
SIM	Service information message.
SNMP	Simple Network Management Protocol.
SVP	Service processor. A notebook computer built into the disk array. The SVP provides a direct interface to the disk array and is used only by the HP service representative.
TID	Target ID.
VSC	Volume Size Configuration is a feature that defines custom volumes (CVS volumes) that are smaller than normal fixed-sized logical disk devices (volumes).
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.

Index